09. AI risk assessment template

> **Template.** Part of the open-source AI App Policy Kit, version 1.0. Replace everything in [square brackets], delete what does not apply, and have the person who runs the business approve it. A starting point, not legal advice.

Business	[Business name]
Document owner	[Name and role]
Approved by, and date	[Name], [date]
Next review	[Date, at most 12 months on]

Why this matters

Before an AI feature touches customers or their data, write down what could go wrong and what you will do about it. Where personal data is involved and the risk to people could be high, UK GDPR requires this as a data protection impact assessment (DPIA), and the ICO expects one for most AI uses of personal data. One page done honestly is worth more than forty pages copied.

When to fill this in

- Before a new AI feature or AI tool goes live.
- When an existing one starts using a new kind of data, a new supplier or a new model.
- When an AI system starts taking actions by itself (sending, booking, paying, deleting).
- Every 12 months for anything still running.

1. What it is

Name of the feature or tool	[ ]
What it does, in two sentences	[ ]
Who uses it	[Staff / customers / the public]
AI supplier and model	[ ]
Assessment by, and date	[Name], [date]

2. The data

Question	Answer
What personal data goes in?	[ ]
Any sensitive data (health, children, finances, identity)?	[ ]
Why do we need it, and what is the lawful basis?	[ ]
Could it work with less data, or with names removed?	[ ]
Where is it processed and stored, and for how long?	[ ]
Does the supplier train on it?	[ ]
Do the people concerned know? Is it in the privacy notice?	[ ]

3. The decisions

Question	Answer
What does the AI decide or recommend?	[ ]
What happens to a person if it is wrong?	[ ]
Does a person check before anything happens? (document 13)	[ ]
Can it take actions by itself? Which ones, with what limits?	[ ]
Could it treat some groups of people worse than others? How would we know?	[ ]
How does a person question or appeal a result?	[ ]

4. The risks

Score likelihood and harm as low, medium or high. Anything left at high after controls needs a director's decision, and for personal data, advice from the ICO before going ahead.

Risk	Likelihood	Harm	What we do about it	Left over	Owner
It states something false and a customer relies on it	[ ]	[ ]	[ ]	[ ]	[ ]
It reveals one person's data to another	[ ]	[ ]	[ ]	[ ]	[ ]
Someone tricks it with crafted text (prompt injection)	[ ]	[ ]	[ ]	[ ]	[ ]
It takes a wrong action (sends, deletes, pays)	[ ]	[ ]	[ ]	[ ]	[ ]
It is unfair to a group of people	[ ]	[ ]	[ ]	[ ]	[ ]
The supplier is down, changes the model or raises the price	[ ]	[ ]	[ ]	[ ]	[ ]
The bill runs away	[ ]	[ ]	[ ]	[ ]	[ ]
[Other]	[ ]	[ ]	[ ]	[ ]	[ ]

5. If you sell outside the UK

- ☐ We checked whether the EU AI Act applies (it can if people in the EU use the system) and which of its risk categories the feature falls in.
- ☐ People are told when they are talking to an AI system.

6. Decision

Go ahead?	[Yes / yes with conditions / no]

Conditions	[ ]
Approved by, and date	[Director's name], [date]
Review date	[ ]

References

- ICO, *Data protection impact assessments* and *AI and data protection risk toolkit*.
- NIST, *AI Risk Management Framework*. ISO/IEC 42001, *AI management systems*.

AI App Policy Kit, version 1.0. © 2026 PugliseWeb LTD, licensed under CC BY 4.0: use it, change it and share it, with credit. The whole kit, and a free Software Health Check of your app against it: https://www.puglieseweb.com/ai-policy-kit